

**HIT Standards Committee  
Privacy & Security Workgroup  
Draft Transcript  
May 14, 2010**

## **Presentation**

**Judy Sparrow – Office of the National Coordinator – Executive Director**

Thank you very much. Good morning, everybody, and welcome to the Privacy and Security Workgroup. Just a reminder that the public has access today only via the phone, but if you wish to follow along with the slide presentation, copies of the slides are available on the ONC Web site for downloading, and that's healthit.gov if you wish to do that. There will be opportunity at the end of the meeting for the public to make comment.

Let me do a roll call here for the workgroup members. Dixie Baker?

**Dixie Baker – Science Applications Intl. Corp. – CTO, Health & Life Sciences**

Here.

**Judy Sparrow – Office of the National Coordinator – Executive Director**

Steve Findlay? Anne Castro?

**Anne Castro – BlueCross BlueShield South Carolina – Chief Design Architect**

Here.

**Judy Sparrow – Office of the National Coordinator – Executive Director**

David McCallie? Gina Perez? Wes Rishel?

**Wes Rishel – Gartner, Inc. – Vice President & Distinguished Analyst**

Here.

**Judy Sparrow – Office of the National Coordinator – Executive Director**

Jim Bialick?

**Jim Bialick – Genetic Alliance – Health Systems Coordinator**

Here.

**Judy Sparrow – Office of the National Coordinator – Executive Director**

Walter Suarez?

**Walter Suarez – Institute HIPAA/HIT Education & Research – Pres. & CEO**

Here.

**Judy Sparrow – Office of the National Coordinator – Executive Director**

Aneesh Chopra? Chris Brancato? John Moehrke?

**John Moehrke – Interoperability & Security, GE – Principal Engineer**

Here.

**Judy Sparrow – Office of the National Coordinator – Executive Director**

Sue McAndrew?

**Adam Green – Progressive Chain Campaign Committee – Cofounder**

Adam Green is here.

**Judy Sparrow – Office of the National Coordinator – Executive Director**

Adam Green, okay, thank you, Adam. I know members of the Privacy & Security Policy Committee were also invited. Paul Egerman, I believe is on the phone?

**Paul Egerman – eScription – CEO**

Yes, good morning.

**Judy Sparrow – Office of the National Coordinator – Executive Director**

Anyone else from the Privacy & Security Policy Workgroup?

**Carl Dvorak – Epic Systems – EVP**

This is Carl Dvorak attending on behalf of Judy Faulkner.

**Judy Faulkner – Epic Systems – Founder**

And this is Judy Faulkner attending on behalf of herself.

**Judy Sparrow – Office of the National Coordinator – Executive Director**

Wow, go Judy. Anybody else in that workgroup?

**Kathleen Connor – Microsoft Health Solutions – Principal Program Manager**

Kathleen Conner.

**Judy Sparrow – Office of the National Coordinator – Executive Director**

Kathleen, thank you.

**Sue Simatan**

This is Sue Simatan attending on behalf of Marianna Bledsoe.

**Judy Sparrow – Office of the National Coordinator – Executive Director**

Great, thank you.

**Beth Morrell**

And Beth Morrell attending on behalf of Carey Shaw from the Children's Partnership.

**Judy Sparrow – Office of the National Coordinator – Executive Director**

Okay, with that I'll turn it over to Dixie.

**Wes Rishel – Gartner, Inc. – Vice President & Distinguished Analyst**

Judy, before you turn it over, can you ask someone for directions on how to make the slides small enough to fit on the screen?

**Judy Sparrow – Office of the National Coordinator – Executive Director**

Dixie, could you help us with that, the slides are very large. Is there a way to reduce that size?

**Dixie Baker – Science Applications Intl. Corp. – CTO, Health & Life Sciences**

It shows perfectly fine for me, so what does it look like on the screen for you.

**Wes Rishel – Gartner, Inc. – Vice President & Distinguished Analyst**

It looks like your screen is bigger than ours.

**Iwana Singeranu**

Yes, it might be a matter of ....

**Dixie Baker – Science Applications Intl. Corp. – CTO, Health & Life Sciences**

I think you can also do view and maximize work space, try that at the top, it's on the blue.

**Wes Rishel – Gartner, Inc. – Vice President & Distinguished Analyst**

View.

**Dixie Baker – Science Applications Intl. Corp. – CTO, Health & Life Sciences**

And select maximize work space.

**Wes Rishel – Gartner, Inc. – Vice President & Distinguished Analyst**

Let's put a grade out here. On my system it says grade out, I can't.

**John Moehrke – Interoperability & Security, GE – Principal Engineer**

It makes it better, but it still is not completely viewable. It's odd that they don't have one that's shapes it to your resolutions.

**Dixie Baker – Science Applications Intl. Corp. – CTO, Health & Life Sciences**

No they don't, this is not as I've mentioned before, this is not near the high quality of what ONC has. It also can crash every now and then, but hopefully it'll hang on. The only other thing I know to do is to change the resolution of the presentation itself, which of course Iwana would have to do, Iwana would have to do herself.

**Iwana Singeranu**

Yes, I have done that, is it making a difference at all?

**John Moehrke – Interoperability & Security, GE – Principal Engineer**

It looks better now.

**Iwana Singeranu**

Yes.

**Dixie Baker – Science Applications Intl. Corp. – CTO, Health & Life Sciences**

Good, good, thank you, Iwana.

**Wes Rishel – Gartner, Inc. – Vice President & Distinguished Analyst**

The other thing you can do is close the left panel and then it looks okay at least on my screen.

**Iwana Singeranu**

Yes, it's now in lower resolution and it's in presentation mode.

**John Moehrke – Interoperability & Security, GE – Principal Engineer**

It's good.

**Iwana Singeranu**

Okay.

**Wes Rishel – Gartner, Inc. – Vice President & Distinguished Analyst**

Yes.

**Dixie Baker – Science Applications Intl. Corp. – CTO, Health & Life Sciences**

Okay, excellent, excellent. I'm really pleased today to have, this is the fourth of our series of educational sessions on various standardization activities around the management of consumer consent permissions. And today I'm very pleased to have Iwana Singeranu, I hope I haven't butchered your name. Iwana is deeply involved in the development of the HL standards around consent management and consent directives, and I'm really pleased to have her here today. Iwana, thank you very much, and I'll turn it over to you.

**Iwana Singeranu**

Thank you for inviting me to speak in front of your committee. My background is in electrical engineering and computer science. It was more or less by accident that I became involved in the software development for clinical information system integration over 16 years ago. And the standards were at that time, as you can imagine, quite mature where it came to semantic interoperability.

In the last 12 years, I became involved in standards development. And I have decided that for the longest time when I was implementing standards I had very little awareness about what standards development organizations were doing. It is possible to have a lot of the doctors that are not exactly involved on a daily basis with the standards to related activities.

In the last 10 years, I've been consulting in software engineering, application architectures, and my focus has been standard-based interoperability. I'm saying this because you will notice throughout my presentation, the focus and the stress is really on implement ability. The ability of these standards to become part of working solutions is very important to me.

The last six years, I've been working with the Department of Veterans Affairs, the Substance Abuse of Mental Health Services Administration, Centers for Medicare and Medicaid Services over the state of Minnesota. And I've become more aware of the issues facing government and payers, as well as large healthcare providers. So that's it in a nutshell. My background and my aspire perhaps as an implementer of standards, who wants to see a solutions come onto the market that makes healthcare more efficient and the patient safety something that we can all be proud of.

The next slide, I'm providing a quick set of milestones summarizing the efforts within HL-7 primarily, but also work that's been done in HIE to specify patient privacy consent and consent directives. One of the earliest efforts to exchange consent directives as scan documents that contains specific metadata summarizing the effective period of a consent who the authenticator and the witnesses are; what the record is covered by the consent. And that discussion was presented very well by my friend John Moehrke a few weeks ago, and I'm sure it's still fresh in your memory. So I'm going to take it from there basically and discuss what's been done since HIE worked on their BPPC standard.

One very important milestone is the development of an HL-7 version 3 specification for the directive consent and that is an approved normative standard in HL-7. And it has been implemented quite widely in Canada. I want to credit Kathleen Connor from Microsoft as the modeling and publishing facilitator for the team that's working on that project. There were a lot of people who were involved, but in any group there is someone who needs to take a more active role to see the standards through, that standard was valued as normative and we did work on a revision of it as drafts for comment in September, 2008.

It was clear, however, that the directive consent of release one was much more focused on Canadian privacy requirements than U.S. privacy or other international requirements. So we started working on an analysis of a consent and privacy concept that may be applicable internationally. The result of that were several draft specifications that resulted into a draft standard for trial use that was completed this year and is now available.

I want to also mention that the SAMHSA and Richard Thoreson provided substantial support from a subject matter expertise in a leadership standpoint. As you know SAMHSA has a great deal of interest in privacy, since they have several regulatory responsibilities in that area. Theirs was a very important subject matter expertise that led to the development of this analysis model. Once we developed the analysis model, we realized that we knew what information would be required to represent the consent directive, and we were looking for an optimal way of representing that consent directive.

And since the clinical document architecture is one of the standards that's mentioned in meaningful use and it's been piloted in several scenarios by the Nationwide Health Information Network, we decided that it would be useful and relevant to have the consent directive information represented using the clinical document architecture released to specification. So we developed the implementation guide, which constrains the clinical document architecture and is intended to represent the structure content of a consent directive.

There are several organizations that have expressed a great deal of interest, especially SAMHSA and the Department of Veterans Affairs. We used the first version of this specification that was valid in January to develop a pilot implementation converting CDA based consent directive into a computer readable access control policy. We actually pick XACML as an example of that. We also created some forms to demonstrate how this type of standard would be operationalized by application developers, just in an interest of making sure that the standard is implementable.

And it is our image of an upcoming milestone, there is an effort underway to harmonize the security. The engineering and execution of a privacy with the business viewpoint in a harmonized security and privacy domain analysis model, and actually the ballot for this standard has just closed. We have comments coming back and Mike Davis is providing leadership in this area. And again, the idea is to combine the business viewpoints, the business requirements, and make sure that there is engineering and security infrastructure support required to bring these ideas to the implementers and to the clinicians.

The next slide, it's a diagram that basically summarizes the consent directive model and the HL-7 implementation guide, as well as upcoming security engineering combined viewpoints as described in the harmonized security and privacy domain analysis. The key idea behind the slide is that once this analysis was completed, and was actually a pretty extensive and long process, we really had a very good handle on what a consent directive should contain; what kind of privacy preferences a patient might express in a given situation, given that the standard is basically a super set of requirements. It is then possible to take this domain analysis model and apply its content to constrain a specific representation and in this case, a CDA. And it could be applied to either purposes, a name or another on the wire representation might come along that would require this type of information.

This is a description of the pilot implementation. And it's again in the interest of demonstrating that the CDA implementation guide is relevant to implementers. We first developed a prototype user interface, an InfoPath Form actually that allowed an end user to enter the information intended for a consent directive and automatically generate a CDA release to conformant and implementation guide conformant XML document.

That XML document in turn could be automatically transformed into a set of access control policies. As we learn from this pilot, what happens really is that the consent directive results in changes to the default policy to the organization sending the information and the organization receiving the information. Sometimes by opening up the disclosure of information from a sender standpoint and imposing specific obligations on the receiver of that data to make it available to specific users or to use it in a specific way or perhaps to disallow the re-disclosure of that information. So that's the overall flow of the implementation pilot.

Next I'd like to revisit a sample workflow that would apply to personal health record management. In the situation where the client would maintain their consent directives based directly on a privacy policy that it is possible for a variety of organizations to request summary data based on a longitudinal medical record; and even have access to protected information including individually identifiable health information in certain cases. But in this situation as you can see, the consent directive rules could be applied automatically. So based on the preferences of the patient, their primary care physician or a specialist could go directly against their personal health record and try to access the information to which they are authorized. And this is actually perhaps a very simple way to visualize how electronic consent directives would be relevant and meaningful in this case.

In other situations though, it would be the patient in the physician's office consenting to a specific disclosure and allowing then perhaps in a case of continuity of care, the transfer of their medical records from one provider to another provider. And in this case, the ability of the sending organization to the code, the consent directive down to a set of access control policies is very important. And it's also important to convey any additional obligations to the receiver of that information.

The next slide actually is elaborating a little bit at the point a consent directive is really in fact a set of obligations that are access controlled policies that would have to be observed by the sender and the receiver of a healthcare information. This is elaborating a little bit on the continuity of care workflow. The patient may allow therefore, disclosure of protected information to another organization. The receiving organization may be under explicit obligations not to re-disclose that. This is actually perhaps one of the most common used cases. It's a little different than perhaps what the initial Nationwide Health Information Network pilots have done, but it's perhaps what we would see a lot in the area of sharing protected information.

Another dimension of enforcing consent directives is the ability of taking a platform independent standard based interoperable representation of a consent or a privacy policy and making that available to security systems to enforce it at the right time. So ideally, starting from a consensus based standard representation, like a CDA release, to an implementation guide or an HL-7 version 3 message, we can instantiate electronic consents and privacy policies that could be exchanged freely between organizations or exchanged based on specific policies; but allowing for these XML documents to use a common terminology would also enable the transformation into executable policy rules.

And as I mentioned earlier, if you read between the lines a consent directive is in fact foreshadowing specific parameters for specific access control rules that are predefined by a policy. In the case of trying

to automate this process, it's clearly important that the way we're expressing the consent electronically and the way we're enforcing it be as close as possible. And also it has to be somewhat consistent with the way the information is represented. This is where meaningful use actually is coming to our aid because it's promoting the use of interoperability standards for exchanging structured semantically encoded information.

The next slide actually is painting kind of a vision of how interoperable standard based policies and consent directives could be exchanged in an automated way to ensure consistency. So presumably a provider would have the ability to receive these consistent representations of both policy and consent directives. And assemble them in cohesive and executable policy rules that would automatically ensure that national state, province, and local jurisdictional policies, as long as consumers own consent directives, could be automatically enforced. And that's really the key operative word here, we're trying to automate this enforcement mechanism to make it as easy as possible to the implementers of clinical information systems and electronic health record systems to adopt the privacy policies.

Another possible use case here would be to allow for automatic notification and publication of new privacy rules. There is an effort underway as you know regarding quality measures, so quality measures can be encoded now in a standard based format, healthcare quality, a measure format specifically. This is actually not unlike the division for moving quality measure communication toward the standard based format. So if national, state, provinces, state or province jurisdictional authorities could use a common representation for privacy policies, they would be able to share that information and also make sure that the privacy rules between jurisdictions are in sync. So that is basically kind of a future vision of policies and ... representation could be used.

And while today I'm speaking more focused on consent directives, I just wanted to let you know that the standard itself also elaborates the content of a privacy policy. Because through our analysis we found it was very important to have a clear idea of why the privacy policies allow us ... patients to protect our information perhaps against abuse. So that's one of the aspects of the standard that I'm not going to cover today since I'm focused on consent directives.

This is actually quickly a UML diagram. A UML class diagram specifically that identifies the key objects that would be part of a consent directive. And that's somewhat color coded for our convenience to distinguish between some of the classes that contain parameters intended to specify various rule parameters. For example, the purpose for the consent, the action that is allowed by consent, whether the action is enabled or disabled, the effective period for a specific consent directive.

And also in pink, the color coded classes represent criteria that would have to be met by the information that is applicable to this consent directive. In the green classes, we're looking at the sender and the receiver of the information and what criteria they have to meet to fulfill the requirements of the consent.

I just wanted to say as much as this diagram is probably not something we want to delve in today, that working at this level of analysis makes it possible to communicate between different communities. And I know from Dixie that you too, your group includes a variety of stakeholders with varied backgrounds. What we found is that in HL-7 we have a similarly broad audience and stakeholder group. And often times getting privacy and policymakers in the same room with engineers requires great care in trying to convey the concepts across these boundaries.

Creating a ... analysis model allowed us to describe the concepts behind a privacy policy or a consent directive in a way that is familiar to the policymakers, but is also clear enough and concise enough that we can communicate it to those who have more of an engineering to privacy. They're looking for ways to

make it operational. In that regard, I think it's a very good tool to bridge the gap between the business viewpoint and the engineering viewpoint and to make sure that there are engineering and security mechanisms in place to make the provisions of a consent directive operational in a electronic health record system.

The classes really that you see here represent objects that would be referenced as part of the privacy preferences of a patient. And the attributes themselves will allow us to reference a specific type of information consumer or a specific type of information. I just wanted to also emphasize that the requirements on which this model were based are actually international. So we have had an addition to U.S. stakeholders, Canadian and Australian stakeholders that were very active.

In this diagram, I'm sort of taking a step up a little bit from the class diagram and identify the various areas that the consent directive analysis are identified. We have consent specifications that refer to the allowing or disallowing purpose for the consent. An effective period and additional conditions that may apply to the receiver of the data. We also have the ability to specify actions based on a hierarchy of operations. And a very important element of this is a reference to a specific privacy policy that this consent directive is constraining in a sense.

Another very important set of attributes and properties that are identified in the model relate to the kind of information that is affected by the consent directive. And again here we're leveraging the efforts in meaningful use and we're assuming that the information would be structured and encoded using standard terminologies. We also assume that there is a relationship between a specific data that is produced through ... process and the diagnosis. There may be a level of sensitivity that is relevant to the consent directive, as I mentioned or I should have mentioned, is sometimes the coverage type, also plays a role in determining whether the specific information element may be disclosed or not.

I just wanted to delve a little bit more in the type of information. In Canada for example, their consent directives refer to categories of information. So you can specify that you wish to disclose or not, laboratory results or a medication. But it's also possible in this model to go to a much more granular level of detail and specify that you wish to disclose a certain result type, so you can go as generic or as specific as necessary. And that basically affords us the ability to support any number of privacy policies and consent directives.

Another element of the consent directive is to specify who the information sender may be, who the information receiver. This is very important for instance in the Social Security Administration case. They wish to receive consent or authorization to gather information from a variety of organizations. They want all the information for the purpose of determining benefits, and again this model will support their needs as well.

In the case of a consent directive, you can go all the way to specifying the type of information receiver to the role of the person who accesses the information and even the identity of that person. And again, I want to emphasize that in any one implementation, you may not need all of these criteria, but they're available; and therefore, by selecting subsets of these criteria you can support a variety of privacy policies. And again, the next slide I'm trying to sort of show you how these various types of parameters are overlaying over the original analysis model that we have constructed.

In a nutshell, the patient or the client of our healthcare system may be able to specify what type of information they wish to disclose perhaps. They can specify if it's diagnosis related, if it's sensitivity related, if it belongs in a specific business area, and the system will use standard terminology to identify that information. Of course, we don't expect the patient to know the codes, but we expect the systems



are going to encode their consent directives to definitely know and understand those codes. There may be a specific purpose for the intended action, so the disclosure may be intended for treatment for payment, for benefit adjudication, the actions allowed or disallowed by the consent specifically identified, and they could be identified at various levels of granularity as well.

What we found initially is that consent directives stay at this higher level of disclosure use collection, while specific security policies will tend to be much more specific and identify an asset operations on the data. The model allows for very specific or very generic specifications.

We can also include within the consent directive, additional restrictions associated with the specific action, and those would apply perhaps to the receiver of the data. The receiver may not be able to re-disclose the data, may have to audit any access to the data once it becomes part of their organizations own record system. And of course, it's very important to identify who may access the data and who may disclose the data.

In the next slide, I just wanted to go to the next step. Now that we have defined whether the consent directive is to specify to meet the needs of privacy policy experts and also security experts, we can then create an implementation guide that would inform how the consent directive may be exchanged using a specific encoding. In this case, the clinical document architecture. This diagram is basically showing you the relationship between the CDA release two specification, basic patient privacy consent, and the consent directive implementation guide for consent directives.

So while the basic patient privacy consent specification is constrained on the scanned document profile developed by HIE, the consent directive implementation guide that I'm discussing is a direct constrain of CDA. But in doing that, it also took into account the need that we have to be able to exchange, scan with signatures, and even to scan documents if necessary. So it does include the requirements that basic privacy consent supports, but it does so in a slightly different way technically.

On the next slide, I have a few examples, and we have used example consent directive forms to validate that the underlying model can support the needs of various stakeholders. What I don't have here is the Social Security Administration's own consent authorization form, which it's very ... with the model and is fully supported by it. But what you're seeing here is various ways in which the organizations allow their patients to request the release of their information or to allow the disclosure of their information for a specific purpose. Either for the purpose of transferring care from one provider to another, and the patients can identify who can receive the data; what data may be received; and even in some cases, does it ask specific obligations on the receiver of the data?

The form I have on the right-hand side is from British Columbia and it's a very straightforward form that simply states that the patient, actually the guardian on behalf of the patient, is agreeing to the disclosure of their laboratory information. So that is the way the information in Canada is looked at from that kind of close range category of information standpoint. In the U.S., our policies for UCFR part two for example, is very focused on the condition that is related with the information. In this case, substance abuse history.

Of course, there are other policies that organizations enforce that have to do with other types of diagnoses, but again, that is the criteria that comes up over and over again. And it's also a very important piece of information for other uses for patient care, patient safety, for accurate billing, so it is important to be able to trace that piece of information to the diagnosis on which it was created.

**Dixie Baker – Science Applications Intl. Corp. – CTO, Health & Life Sciences**

Iwana, this is Dixie, that slide was probably a good one for those who can't see the webinar, they probably were able to sync up with the slide presentation on that one. But it would be useful for you to say which slide you're on when you go, next slide, just briefly, because I think a lot of our people may have only the slides.

### **Iwana Singeranu**

Yes, I'm sorry, I should be more aware of that. So we're moving from the paper based examples of consent directives to a consent directive form, which is basically an attempt to see that of the consent directive implementation guide, is it implementable by an information system? In this case, instead of filling out a paper form, someone would be able to fill out an electronic form; and in the process of doing that, populating a CDA document. And that's really all that this slide is trying to show by filling in the fields in the form. We're effectively populating a specific element and attributes in the clinical document that you can see on the right-hand side. While this first consent directive form slide is focusing on a lot of the header information that identifies the patient whose record is affected, the organization that holds the information, and the document of the organization that is intended to receive the document.

The next slide is looking at some of the consent details that would be part of the body of the document and you may specify a purpose of use. For instance in this particular form, you can enable or disable the disclosure and you can identify the types of information that would be relevant. In this case as you can see, the information is referring to specific documents. So the granularity of the disclosure is entirely up to the policymakers really and it's supported by the standards.

On the next slide, we're going to take another look at the cryptic XML, just to say that in creating a consent form we're explicitly creating a CDA document that has all the structure and all the details of the consent directives domain analysis model.

The next slide, we're looking at an electronic consent directive form as a rendered CDA document from an HDML and also in its original form. And I just wanted to say that basically the CDA document itself has a document ID. It could be referenced and it could be used to specify how a specific information artifact belonging to this patient may be handled. So you can relate now a continuity of care documents let's say with a consent directive that applies to it. For as long as this consent directive is valid, those provisions would apply to how the content of the continuity of care document would be handled by someone who let's say is pulling this information off a registry into a National Health Information Network. It is possible now to match up in a way that is aware of the privacy concerns of patients and allows you to specify the consent directives applicable to certain information artifacts and clinical documents.

In the next slide, we're looking at the details actually. Simply the purpose of the slide is basically to allow you to recognize the various criteria and different attributes that were elaborated in the consent directive domain analysis model, and show you that they actually appear in the form document. It took us a little bit of time and a little bit of analysis to see that each one of these forms is actually talking about specific types of information.

There's information about the action and whether the action is enabled or not. There's information right there in that narrative about who the sender and the receiver of the information is. And explicit information about the types of information that may be disclosed and what obligations would have to be met by the receiver of that data. It is actually again as a sanity check, as well as the source of requirements for us. These forms were very useful because they basically tell us what are the elements of a consent directive.

And then the next slide, I just wanted to basically focus on the fact that since CDA is the underlying format for continuity of care documents, it is possible for an organization to render and look at the continuity of a consent directive in the same way that they would treat a continuity of care document. In terms of information exchange and rendering of the content, it makes it very easy for adoption, that's one of the reasons why we selected this encoding mechanism.

In summary, the pilot project showed us that we can take a form, we can create a CDA document by populating the fields in the form, and what we're creating as a result of that is an interoperable representation of a consent directive. And again, I would like to just leave you with a thought that structure is very complicated. And the ability to create it automatically and focus in on those classes of objects and those attributes are very important for downstream.

Processing of these consent directives were foremost in our minds. And that made it possible to then go back and abstract again the structure when we created the form. The form itself focuses only on those lead elements that are relevant from a business standpoint, but there is a lot more information that is encoded in the body of a CDA document that is not immediately irrelevant or exposed in the form.

The next slide, just wanted to make a few points about what happens when you're trying to apply the consent directive to create a continuity of care records or documents. We're very much dependent on specific quality of data. So we have this very nicely structured consent directives that are using better terminology that ... specific attributes of the data to determine whether it's protected or not. But if the data itself has not ... that level of quality we certainly cannot automate the enforcement of these policies. By the way, can you hear me?

**Dixie Baker – Science Applications Intl. Corp. – CTO, Health & Life Sciences**

Yes, fine.

**Wes Rishel – Gartner, Inc. – Vice President & Distinguished Analyst**

Yes.

**Iwana Singeranu**

I have this sound in my earpiece. Basically, I just wanted to say that we're at a very propitious time right now where we're in a position to automate a lot of these privacy concerns simply because the quality of the data is improving as well. So the same requirements derive from meaningful use that lead to standardizing the structure and the encoding of the data by electronic health record systems is also enabling better decision support tooling that will lead to better patient safety. It also enables accurate billing, because now the billing items can be directly associated with the events and the documentation that's captured by the electronic health record system.

It's certainly the same type of drivers and requirements that allow us to itemize data for public health reporting. So all these different requirements come together to improve the quality of the medical records, the electronic health records themselves, and allows us to better protect that, which is deemed sensitive by a policy. We can create rules that automatically enforce the separation perhaps of protected data from non-protected data, and it's done basically as part of the workflow. It actually is a very good time to have this kind of discussion, because we can see the data is getting to the point where these policies could be automatically enforced.

And it is possible also to put the patient's privacy concerns at rest, because we can explicitly reference a consent directive from a document or a message. If the information is expressed electronically, the consent directives are expressed electronically in a standard based interoperable platform mutual way,

and we have the ability now to relay these two together. And we don't have to pick winners and losers in terms of technology, but we can allow the industry to make the best possible choices.

And the next slide, actually I have an example of how we could reference a consent directive from a continuity of care document. You can see in the CDA structure, the confidentiality code is one of the elements of the header of the document, but it's also present in each one of the sections of the continuity of care document. We could reference that whole document covered by a specific consent directive or we could even go down to a specific section of a document and specify that. And again, this is specifically supported by standards right now that are a part of the meaningful use criteria.

This is the next step now in our pilot implementation. We were able to describe what a consent directive contains. We were able to encode that in CDA. We were able to automate that, the creation of those CDA based consent directives. And then we wanted to see if we could actually leverage that to transform what is basically a platform mutual representation into something that could be directly processed by a rules engines.

And XACML is also a standard produced by OASIS. And there are several implementations of the standard that would process these rule sets accordingly. So we thought that would be a good example of again one possible implementation of privacy that would leverage XACML. And what I wanted to show you here is how a specific set of criteria in our consent directive for instance, the effective time of the consent directive becomes a rule, part of a policy set in XACML. And a very interestingly thing is that our documents architecture representation allows us to select the information and then reference it so we can reference the expiration and a start date of a specific policy based on their express path location in the CDA document.

So there is a unique identifier based on the structure of a CDA document that will reference the expiration and the start date for a consent directive. I thought that was a particularly useful way in which we can identify the parameters of a new rule that says while this consent directive is in effect, permit access or permit the action that is associated with it. So this is just part of an entire consent directive, but it's the part that addresses the effective period.

Another example is, how does the intended recipient of a consent directive become an XACML policy? And this is a rule that deals with the recipient name being allowed by the consent directive, so we have the intended recipient ID. There are two different attributes that are part of that, the root and the extension, and that is identifying the allowed recipient of that information. And again, the wonderful thing about this, and it's probably hard to get as excited as I am about it, is that I have the discriminator that tells me exactly what information I'm referring to that's based on my CDA encoding of a consent directive.

I have a couple of resources here that I wanted to make you aware of. The work is still ongoing, and while we're very close to completing it, I just wanted you to be aware it's part of the resources slide. Another thing I wanted to mention is the receiver obligation, we've already mentioned that, and XACML has supported this and tactically supports these obligations as part of a rule. And again, we're able to leverage the structure of this CDA implementation guide to identify for instance, how long or what event has to occur before a specific type of information is no longer available to the receiver.

And this is my last slide with additional references for you. I wanted to let you know that the draft standard for trial use allows you to enter comments. This gives us an opportunity to gather feedback from people who are trying to apply this analysis model to their purposes. And that the implementation guide is undergoing a ballot of the cycle, and if everything goes well, it should be approved later in June/July timeframe. That was all I had prepared.

**Dixie Baker – Science Applications Intl. Corp. – CTO, Health & Life Sciences**

Thank you very much, Iwana, this was really, really good. So let's go into a period where our workgroup can ask Iwana questions.

**Dave Wanser – NDIIC – Executive Director**

This is Dave Wanser. One of the questions, I've got two questions actually, one is you said this is actually applicable and useable now. As we look at making policy, what is a reasonable kind of trajectory for this being something that you would expect would be an absolutely universal use?

And then the second question is, given the differences in state laws around confidentiality, does it not get very, very complicated to try to specify all those nuances in this framework?

**Iwana Singeranu**

I think what you're referring to is that any consent directive or any privacy policy represents a set permutations of these different criteria. What we tried to do was to find those criteria that could be subject to permutation. So in so far that all the policies issued by state deals with the criteria that are defined in detail in this model, then this model will support any policy.

If they're out of criteria that we haven't thought of, then really that's always—So from a policy expert standpoint, I would expect that you would look at this model and you'd say, you didn't think about these particularly attributes of an information artifact that would make it protected or you didn't think about this particular obligation that a receiver of information ... Because without that we cannot deploy this type of automation policy assessment in the state of Oregon let's say.

Basically, what we tried to do was not come up with all the possible policies that may exist out there, that's really your job, but we wanted to give you the language so you can declare or define any policies that might be in scope for let's say the United States. It's a little different, it's providing the alphabet really rather than writing the dictionary. The way we see it as that the policymakers are writing the dictionaries. And all we have done here is to create the language by which you can support all the various words that you can think of.

**John Moehrke – Interoperability & Security, GE – Principal Engineer**

Iwana, this is John Moehrke. I'd also add that this is a palliative process ongoing into DSTU. The intention going into DSTU is to allow for experimentations as to whether this meets needs or not.

**Iwana Singeranu**

Exactly.

**John Moehrke – Interoperability & Security, GE – Principal Engineer**

This could very well be policies that are discovered during this draft standard for trial use phase that might affect the standard as it goes further into maturity.

**Iwana Singeranu**

Absolutely, yes. And that's exactly what I've done. We need that kind of feedback back from people who are trying to apply it and then to say, there is a criteria in here that's missing. We need this criteria to express this consent directive or this policy. Again, I did not delve into the policy model, which is very similar, but that's exactly the mechanism that we like to engage the stakeholders to tell us what else would have to be available. Again, to complete the alphabet, to have all of the letters you need, but you can assemble that in the nicest way.

**Kathleen Connor – Microsoft Health Solutions – Principal Program Manager**

Iwana, this is Kathleen. I would also emphasize that the methodology that you're using allows for extensions without having to change the standard and also allows for the addition of co-systems and terminologies without to actually specify policies without having to change the model.

**Iwana Singeranu**

That's true, yes, especially on the terminology end of things; especially in this model we're basically identifying the concepts and allowing the terms to be specified perhaps on a project-by-project or nation-by-nation basis. This is actually an international model, could be localized appropriately.

**Wes Rishel – Gartner, Inc. – Vice President & Distinguished Analyst**

Iwana, this is Wes. I have a couple of questions. I'd like to comment that every standard I've ever worked in has a sense of feeling done until it's implemented and then we redefine done by our experience in implementation. I'm delighted that you have the draft standard for trial use process available to advance this. And wondering if it would be possible, as well as collecting comments on the contents of the standard, collect comments on the implementation success or best practices for implementing during the time period.

Because I think that a lot of what it takes to get to standards on a regional or national level is an understanding of the impact on systems that control the release of data and that have to categorize data for release. And then to implant in organizations of making what has been at best verbally written, and therefore ambiguous policies to be actually unambiguous. I noticed a lot of codes in the DAM, the model.

**Iwana Singeranu**

Hello?

**Dixie Baker – Science Applications Intl. Corp. – CTO, Health & Life Sciences**

Hello?

**Iwana Singeranu**

Hello?

**John Moehrke – Interoperability & Security, GE – Principal Engineer**

Yes, Wes, did—

**Iwana Singeranu**

Did we lose Wes? I thought, okay. Well actually, I just wanted to say quickly that he probably wanted to point out that there are a lot of coded attributes in the domain analysis model that would rely on the standardized terminology. We do have obligation codes, rule codes, basically object codes, and diagnosis codes.

And again I think where the opportunity lies is as these codes are identified for the purpose of meaningful use for other purposes having to do with improving the overall quality of healthcare, they can also be reused to represent the consent directives and the privacy policies. And therefore, the privacy policy and the consent directives would be directly supported by the underlying data.

I don't know if Wes has rejoined us?

**John Moehrke – Interoperability & Security, GE – Principal Engineer**

This is John Moehrke. I'll add onto that, I think two other things, one is this audience may not understand that the domain analysis model is a high-level view and where we've actually put in specific attributes, it's primarily informational. When looking at the domain analysis model, that may or may not be the concrete vocabulary use, but probably is.

But it's also, the other piece that we've struggled with, if an attribute is not there that you think should be, that's actually pretty typical, because the domain analysis model is just trying to show the relationships of classes. And the content of attributes within a class is there to kind of help you understand what that class means, not to say it normatively has to include those attributes.

**Iwana Singeranu**

Right. We did try actually as much as we could though, John—

**John Moehrke – Interoperability & Security, GE – Principal Engineer**

Yes.

**Iwana Singeranu**

—identify all those that we saw. So if there are any that are important and don't appear in the domain analysis model, then that's exactly what we need to hear from the implementation community.

**John Moehrke – Interoperability & Security, GE – Principal Engineer**

The other thing I wanted to address to the specific, Wes' comment, was that there is another project ongoing within HL-7 to create a catalog of well-known privacy policies. The idea of creating a privacy policy that could be reused if a particular region agrees that that policy meets their needs, it could then be given a well-known policy identifier. So there is a work item in HL-7 to create a catalog of well-known privacy policies and what they would mean. That was another thing he was alluding to.

**Dixie Baker – Science Applications Intl. Corp. – CTO, Health & Life Sciences**

Yes, this is Dixie, my—

**Wes Rishel – Gartner, Inc. – Vice President & Distinguished Analyst**

Yes, this is Wes, I just got disconnected just as I got to the end of my questions there. I don't want to occupy the group with re-answering, but maybe Iwana, you can send me an e-mail offline or something to let me know.

I think what I'm driving towards is to look at this policy and identify those steps such as reviewing the states, identifying, getting some real world experience, and creating the catalog of all the artifacts necessary to roll this out at scale.

**Iwana Singeranu**

One thing that John was mentioning is there a work item, actually the community based collaborative of care and the security workgroups in HL-7, to look at identifying basically well-known privacy policies and identifying them easily. And that would be an effort basically to exercise this domain and the CDA implementation guide.

**Walter Suarez – Institute HIPAA/HIT Education & Research – Pres. & CEO**

John and Iwana, this is Walter. I think it will be helpful to distinguish or to clarify the term policies in the context of HL-7 projects with the policies in the context of federal and state laws and regulations. There's a lot of tendency to confuse perhaps that when in the HL-7 realm we talk about policies. We're not necessarily talking about, not covering necessarily—

**Dixie Baker – Science Applications Intl. Corp. – CTO, Health & Life Sciences**

Walter?

**Walter Suarez – Institute HIPAA/HIT Education & Research – Pres. & CEO**

—several laws or regulations. John, if you could perhaps or Iwana, address that, might that be helpful?

**Dixie Baker – Science Applications Intl. Corp. – CTO, Health & Life Sciences**

Can we first have whoever has all the background noise, please mute your phone. Thank you very much, okay.

**Iwana Singeranu**

Well, the first thing I'd like to say is that in the analysis of the consent directive and the privacy policy we did not attempt to inventory every privacy policy out there. We tried to identify what criteria they have in common and what criteria would be required to represent both the consent directive and the privacy policy. And that's what we captured in this model.

What John was referring to is an effort that is still in its very early stages to try to inventory explicit privacy policies. And I don't have any other information except that is a future work item.

**Wes Rishel – Gartner, Inc. – Vice President & Distinguished Analyst**

My suggestion is something that this committee might consider recommending up through its channel, is to conjoin that effort or to get explicit cooperation between that effort; and other work that ONC and the policy committee have going on to identify issues around implementations of HIEs at the state level. Because I think there has been a lot of good work done on both sides. If we can bring that work together and create that particular artifact that it otherwise becomes a part of each implementation project, we can go a long way towards proving that the analyzed structure is in fact adequate; and streamlining implementations as this turns into required standards and gets rolled out.

**John Moehrke – Interoperability & Security, GE – Principal Engineer**

I very much agree with you, Wes. Being able to have proof that something works are better than having simply assertions that it works.

**Wes Rishel – Gartner, Inc. – Vice President & Distinguished Analyst**

Well, yes, I think that's an obvious requirement. But I think we have an opportunity here because of the breath of different activities that are going on right now and during health IT to anticipate that issue. And that's why I'm suggesting that this workgroup create a recommendation, send it up through the committee, and to ONC.

**Dixie Baker – Science Applications Intl. Corp. – CTO, Health & Life Sciences**

Yes, I agree with you, Wes. And I would further add, as I heard this and as I've had some recent conversations with some people from the VA who are implementing some of this as Iwana mentioned, I think that there's no question in my mind that you can implement this on a computer, no question.

But I think that what we really need to look at from the charter of our committee in particular are twofold. Number one is, how to implement it in a way that the consent directives are truly interoperable between organizations, and that touches on your comment about the HIEs, Wes. But secondly, something Wes I think you were touching on, and that is not only the technical implement ability, but the operational implement ability. Just because a computer comes up with an answer doesn't necessarily mean that it



really can be implemented on the operational level. And I think that our recommendation should address both of those.

**David McCallie – Cerner Corporation – Vice President of Medical Informatics**

Yes, this is David.

**Iwana Singeranu**

Hello.

**David McCallie – Cerner Corporation – Vice President of Medical Informatics**

Can you hear me?

**Dixie Baker – Science Applications Intl. Corp. – CTO, Health & Life Sciences**

Yes.

**David McCallie – Cerner Corporation – Vice President of Medical Informatics**

Good, sorry. I'm worried about the operationalization of it, becoming a little bit of a skeptic, a skeptical point of view. I think on the good side it's tremendous work for and very valuable for sort of categorizing the intent of a consent policy capturing in a structured fashion the parameters that might govern someone's intent.

Mapping that intent to actual codes, in an actual system where you have overlapping policies that may be in conflict, that may have overlapping expiration dates, it is incredibly difficult, even the simplest things. We've been involved with a local HIE here in Kansas City, and just this simple question, can you filter out sensitive drugs and diagnoses so that we don't have to worry about managing the state-by-state variations in this Kansas City area where we span two or three states? It was incredibly difficult to define, what is a sensitive drug, what is a sensitive diagnosis, what is a sensitive lab result, what is a sensitive combination of lab results, which on their own might not be sensitive? So the operationalization of this in a computable form is still extremely difficult.

And I'm curious to know what, if anything has been done on the nomenclature side of this to try to define things like sensitive drugs or sensitive labs or sensitive diagnoses. In a sense it's kind of a pointless question, but in the real world it's actually the real question. I'm curious to know if anything has been done to try to create kind of a taxonomy of how these would actually map to real world codes.

**Iwana Singeranu**

I think that's actually a very good question, because that's precisely what the analysis tried to find out. What I did different, attributes of the data. So you said sensitive drugs. Well aspirin is not sensitive, but if it's associated with HIV, an HIV diagnosis, is it sensitive or not? I don't know. AZT may be is sensitive by itself, because it's an indicator of another diagnosis. But the valuable aspect of it is that you can, by policy or by consent directive, associate the different elements that are sensitive. So maybe AZT on its own is not sensitive, maybe it is, maybe aspirin itself is not sensitive, maybe it is, but maybe the result of specific test is what makes it sensitive. Again all of those things are already controlled by standardized terminology, so it is up to the policy to bring them all together.

**David McCallie – Cerner Corporation – Vice President of Medical Informatics**

Yes, that's well and good, but that doesn't take us any further than what the paper took us. Those rules could have been expressed on a piece of paper. If you want to turn it into something where an algorithm can actually make a decision and filter something, you actually have to get concrete like that.

**Iwana Singeranu**

Right.

**David McCallie – Cerner Corporation – Vice President of Medical Informatics**

And consumers can't do that. They don't know enough.

**Iwana Singeranu**

That's right.

**David McCallie – Cerner Corporation – Vice President of Medical Informatics**

Even our expert group of physicians when we presented some of the choices to them and said, should this be considered a sensitive drug or not? They kind of shrugged their shoulders and said, it depends, and that kept coming up so many times of it depends.

**Iwana Singeranu**

And all of the time it depends on the policy too. Again, let me just emphasize what I was trying to say earlier, that, yes, if the policy is encoded beautifully, but the data is not, then the policy cannot be automatically applied to the data.

**David McCallie – Cerner Corporation – Vice President of Medical Informatics**

We have fully coded data, and let's say we have fully coded policy, the problem is the mapping between those two.

**John Moehrke – Interoperability & Security, GE – Principal Engineer**

Yes, but your point is absolutely true. And no, this did not create a nirvana that somehow clarifies that particular problem. What this is trying to do is take a next step beyond a simple assertion consent, which BPPC is, to the ability to have embedded into the patients specific consent document some additional attributes that are applied to the policy, but it did not solve this unsolvable piece, which you're bringing up. It's still very true.

And I think this is related to the question that Walter asked, and certainly the policy writing is a very difficult thing to do. And it does start from looking at federal policy, looking at state policy, looking at ethical policy, looking at professional society policy, and ultimately all that this is trying to represent is that fragment of the overall privacy policies that are specialized for this specified patient.

**Iwana Singeranu**

Right. This is—

**John Moehrke – Interoperability & Security, GE – Principal Engineer**

But also a whole lot of privacy policies that is outside this particular object. And I think Walter was probably also trying to tee up a reminder that in HITSP we wrote this PN900 that also talked about all of the policy building as being an exercise of policy writing that is if you don't have that. It doesn't help you at all to have a very expressive consent mechanism if it doesn't have that overall policy framework that it fits within.

**David McCallie – Cerner Corporation – Vice President of Medical Informatics**

This is David again, and I certainly agree. I don't want to be taken as criticizing this work, I think this is brilliant work and extremely useful. I just want to caution those downstream from us who actually have to put it into practice in say for example an unattended filtering model in an HIE; that even if you have all of

this codified policy actually mapping it to real world decisions is incredibly difficult to the point of where we ended up almost giving up, it was so hard.

**Iwana Singeranu**

This is just—

**David McCallie – Cerner Corporation – Vice President of Medical Informatics**

And this doesn't change that, it just makes it cleaner what our next set of problems are.

**Judy Faulkner – Epic Systems - Founder**

This is Judy and I agree with that. And we've had trouble with folks in a number of different areas where they've talked about exactly that, how do they segment that? And we've asked them, can you tell us what you can and can't send through? And they cannot do that, they just say it's not doable.

I have a big concern about what we reviewed and that is from a computer science point of view, it is a good flow chart representation of how to do privacy. But it isn't going to work until you map it together with how to do care, because the two can't be separated, and right now it's alone. And if in fact you throw in how do you do care with it, I think it will change drastically.

And then I think you also, not only will it change drastically, because underneath it all we're trying to take good care of patients. And if in fact we can't do both in one by separating them, we have to put them together and see if it will work together. I think that the question of data segmentation is still a decision that hasn't been made. So when I see this going forward with the assumption that there is data segmentation, I'm confused about the process where I thought that that is not something that has been decided yet.

**Kathleen Connor – Microsoft Health Solutions – Principal Program Manager**

This is Kathleen.

**Wes Rishel – Gartner, Inc. – Vice President & Distinguished Analyst**

Yes, I think we all are having the experience that was by even a Larson Cartoon some years ago. Where there's a blackboard full of equations, and then right in the middle of the flow of equations here a miracle occurs, and then the equations pick up again. And I see the draft standard for trial use, various implementation tests, as a way of increasing awareness, as David said, where the next issue is. And then we're either smart enough to find a way to finesse that or we recognize where the real challenge is and we don't attempt to roll something out nationally that is going to hit that stumbling point.

**Judy Faulkner – Epic Systems – Founder**

I just—

**John Moehrke – Interoperability & Security, GE – Principal Engineer**

... this is—

**Judy Faulkner – Epic Systems – Founder**

Can I just—

**John Moehrke – Interoperability & Security, GE – Principal Engineer**

... is the capability ...

**Dixie Baker – Science Applications Intl. Corp. – CTO, Health & Life Sciences**

John, John? I think we need to bring this to a close, I think that's what Judy is trying to say.

**Judy Faulkner – Epic Systems - Founder**

I wasn't trying to say bring it to close, I was—

**Dixie Baker – Science Applications Intl. Corp. – CTO, Health & Life Sciences**

I thought it was Judy Sparrow, I thought it was—

**Judy Faulkner – Epic Systems - Founder**

Okay, no, this is Judy Faulkner.

**Dixie Baker – Science Applications Intl. Corp. – CTO, Health & Life Sciences**

Okay, okay. I—

**Judy Faulkner – Epic Systems - Founder**

No, I was saying that without meshing it with care, it's an incomplete process.

**John Moehrke – Interoperability & Security, GE – Principal Engineer**

Yes, but this is not a process, that's what I'm trying to point out. This is only representing a standard for how the standard has a capability to do, not that you must do it this way.

**Wes Rishel – Gartner, Inc. – Vice President & Distinguished Analyst**

I think everybody is giving full credit to the standard assuming it's been well analyzed and well implemented and then asking the question, so now what do we do with it?

**Iwana Singeranu**

Well actually, can I just—

**John Moehrke – Interoperability & Security, GE – Principal Engineer**

And that's very well taken and that's where I would see to, this has to be put into a care setting ...

**Dixie Baker – Science Applications Intl. Corp. – CTO, Health & Life Sciences**

Iwana is trying to make a comment, and I do think, let's let Iwana make her comment and respond to that, but then we do need to bring this meeting to a close, we're already 20 minutes over time.

**Iwana Singeranu**

I just wanted to make a quick point that, as I mentioned already, there is a privacy policy analysis as well in this DSTU. Today I focused on a consent directive because that was the focus I understand for your workgroup. But yes, there are two points I'd like make, that we're looking for similar feedback to make sure that from the standpoint of the person writing the policy that there are all the different criteria required to write the policy. Now I agree that writing the policy itself is not a trivial matter. And determining what is protected and what is not protected is not a trivial matter.

And I just wanted to go back to John's point, the standard makes it possible for you to encode it in such a way that it becomes interoperable and could be translated into an access control ... management language, that's really the scope of it. And it's not really dealing with the workflow as much as it allows you to automate the exchange of a consent or a privacy policy in ... space.

The healthcare itself, the healthcare delivery process, it's still its own process, and this is just an ... to make it possible to capture the privacy preferences and represent them in a standard based way. That's all I had to say.

**Dixie Baker – Science Applications Intl. Corp. – CTO, Health & Life Sciences**

Okay, thank you very much. I think we better open the lines up for public comment, Judy?

**Judy Sparrow – Office of the National Coordinator – Executive Director**

Okay, that will be great. Operator, could you open the line for any public comment please? Hello?

**W**

Melissa?

**Moderator**

If you would like to ask a question, please press star one on your telephone keypad at this time. Our first question is from Jim Kretz with the Substance Abuse and Mental Health Service Administration, please proceed with your question.

**Jim Kretz – SAMHSA – Project Officer**

I just wanted to comment on this last bit of the discussion that finding a particular set of drugs that are somehow sensitive. I believe that's an absolutely useless quest that you have to rely on what the ... purpose, because their sensitivity is what's an issue.

**Judy Sparrow – Office of the National Coordinator – Executive Director**

Okay, thank you, Mr. Kretz, anybody else on the line?

**Moderator**

There are no other questions at this time.

**Judy Sparrow – Office of the National Coordinator – Executive Director**

Okay, thank you. Dixie, do you want to make some final closing remarks?

**Dixie Baker – Science Applications Intl. Corp. – CTO, Health & Life Sciences**

Sure. I want to thank once again Iwana for an obviously very interesting and well thought through and very valuable presentation on a very admirable work. And you could tell by our questions and the fact that we went way over time, we have a lot of interest in this area, and we'll continue to watch as DSTU continues to move forward.

I'd also like to thank those of you who dialed in. It's really important that our decisions that we make both on the Privacy and Security Policy Workgroup and the Standards Workgroup are well informed and these series of educational sessions are aimed to help us do that. So thank you all.

**Judy Sparrow – Office of the National Coordinator – Executive Director**

Thank you, Dixie.